

8 Scepticism

8.1 PLAN

Rational thinkers respect their evidence. Properly understood, that is a platitude. But how can one respect one's evidence unless one knows what it is? So must not rational thinkers know what their evidence is? If so, then for rational subjects the condition that one has such-and-such evidence should be non-trivial yet luminous. But how can it be, given the anti-luminosity argument of section 4.3?

The assumption that rational thinkers know (or are in a position to know) what their evidence is has implications for sceptical arguments. Non-sceptics postulate a special asymmetry between the good and bad cases in a sceptical argument (section 8.2). Sceptics try to undermine the asymmetry by claiming that the subject has exactly the same evidence in the two cases, but this claim is not obvious (section 8.3). We can argue from the premise that rational thinkers know what their evidence is to the conclusion that their evidence is the same in the two cases (section 8.4). That conclusion forces one into a phenomenal conception of evidence (section 8.5). But the premise that rational thinkers know what their evidence is leads by a parallel argument to a clearly false conclusion (section 8.6). This is another variation on the arguments of sections 4.3 and 5.1. Rational thinkers are not always in a position to know what their evidence is; they are not always in a position to know what rationality requires of them (section 8.7). These conclusions generalize to sceptical arguments in which the sceptic does not claim sameness of evidence between the good and bad cases (section 8.8). One upshot is that sceptical arguments may go wrong by assuming too *much* knowledge; by sacrificing something in self-knowledge to the sceptic, we stand to gain far more in knowledge of the world.

8.2 SCEPTICISM AND THE NON-SYMMETRY OF EPISTEMIC ACCESSIBILITY

For simplicity, we can treat the sceptic as a generic figure, without attempting to track the protean variety of sceptical argument. Scep-

Scepticism

165

ticism is a disease individuated by its symptoms (such as immoderate protestations of ignorance); we should therefore not assume that it can be caused in only one way. The present aim is to identify one main such way, not to eliminate the disease entirely.

For the sake of argument, let us assume that the constraints of the content externalism discussed in sections 2.2 and 3.2 are consistent with grasping the relevant propositions in sceptical scenarios. A recently envatted brain can still think about the external world. Even for such a brain, the assumption remains problematic as applied to propositions expressed by means of perceptual demonstratives (see also section 7.6). Suppose, for example, that I am looking at a cloud and think that that cloud is dark. A brain envatted before the advent of that cloud, with experience in some sense indistinguishable from mine, does not think that *that* cloud is dark, although it may think the words 'That cloud is dark' and be in no position to know that it does not thereby express a singular proposition concerning some cloud to the effect that it is dark. Similar issues arise for much less extravagant sceptical scenarios, involving mere hallucinations and the like. We assume for the sake of argument, perhaps over-generously, that the sceptic has some way of absorbing such implications of content externalism.

The sceptic compares a good case with a bad one. In the good case, things appear generally as they ordinarily do, and are that way; one believes some proposition *p* (for example, that one has hands), and *p* is true; by ordinary standards, one knows *p*. In the bad case, things still appear generally as they ordinarily do, but are some other way; one still believes *p*, but *p* is false; by any standards, one fails to know *p*, for only true propositions are known. As far as externalism permits, things appear to one in exactly the same way in the good and bad cases. The sceptic argues that because one believes *p* falsely in the bad case, one does not know *p* (even though *p* is true) in the good case. Let us postpone asking why the sceptic should think that false belief in one case precludes knowledge in the other, and consider the bad case.

Uncontroversially, if one is in the bad case then one does not know that one is not in the good case. Even if one pessimistically believes that one is not in the good case, one's true belief does not constitute knowledge; one has no reason to suppose that the appearances are misleading to that extent. More generally, it is consistent with everything one knows in the bad case that one is in the good case. For even if in the bad case one believes some true propositions which entail that (contrary to the appearances) one is not in the good case, those true beliefs do not all constitute knowledge. Part of the badness of the bad case is that one cannot know just how bad one's case is.

For the sceptic, the two cases are symmetrical: just as it is consistent

with everything one knows in the bad case that one is in the good case, so it is consistent with everything one knows in the good case that one is in the bad case. One simply cannot tell which case one is in. For the sceptic's opponent, the two cases are not symmetrical: although it is consistent with everything one knows in the bad case that one is in the good case, it is not consistent with everything one knows in the good case that one is in the bad case. For in the good case, according to the sceptic's opponent, one knows p (for example, that one has hands), and also (by description of the bad case) that if one is in the bad case then p is false. These three propositions are jointly inconsistent:

- (a) One is in the bad case.
- (b) If one is in the bad case then p is false.
- (c) p .

That argument does not assume that one knows that which is a logical consequence of what one knows, for the anti-sceptic's conclusion was merely that it is inconsistent with what one knows in the good case that one is in the bad case, not that one knows in the good case that one is not in the bad case. Although the anti-sceptic may hold that in the good case one also knows that one is not in the bad case, the asymmetry does not require that further knowledge claim.

We can state the asymmetry in the terminology of epistemic logic (see also section 10.4). A case β is said to be *epistemically accessible* from a case α if and only if everything which one knows in α is true in β . Then, according to the anti-sceptic, although the good case is epistemically accessible from the bad case, the bad case is not epistemically accessible from the good case.

Some refinements may be needed to handle the issues raised by the broad content of indexical expressions. As uttered in any case α , the sentence 'This case obtains' expresses a content true in α and in no other case. Perhaps one can know that content in α without knowing everything about α ; we might allow cases other than α to be epistemically accessible from α , on the grounds that 'This case obtains' expresses (different) true contents in them. This complication does not affect the main arguments to come.

As is well known, asymmetries of epistemic accessibility yield counterexamples to the epistemic version of the 'Brouwersche' thesis in modal logic, the principle that if p is false then one knows that one does not know p ($\neg p \supset K\neg Kp$; 'K' for 'one knows that'), and consequently to the epistemic version of the S_5 thesis, the principle that if one does not know p then one knows that one does not know p ($\neg Kp \supset K\neg Kp$).

The latter principle entails the former because knowledge is factive ($Kp \supset p$ holds). Like epistemic S_4 (the KK principle), epistemic S_5 embodies a luminosity claim. But the failure of the epistemic S_5 principle on non-sceptical assumptions was already noted in section 1.2, independently of general anti-luminosity arguments. For in the bad case, p is false and one does not know p , but one does not know that one does not know p . If one knew in the bad case that one did not know p , then according to the sceptic's opponent it would not be consistent with everything one knew in the bad case that one was in the good case, since these three propositions are jointly inconsistent:

- (d) One is in the good case.
- (e) If one is in the good case then one knows p .
- (f) One does not know p .

According to the sceptic's opponent, one can know (e) even in the bad case by description of the good case and one's appreciation that it meets the conditions for one to know p . The failure to know that one fails to know is characteristic of the bad case. Although the sceptic will try to argue that the postulated asymmetry between the two cases is ultimately unstable, there is at least no immediate incoherence.¹

A common means of slurring over the epistemic asymmetry is to speak of the two cases as *indiscernible*. Surely, if x is indiscernible from y then y is indiscernible from x . But even indiscernibility embodies a concealed asymmetry. For one may be able to discriminate between x and y when they are presented in one way and not when they are presented in another (Williamson 1990a: 14–20). A case can be presented in two relevant ways. When one is in a case, one can present it *indexically* to oneself, as 'my present case'. Alternatively, whether one is in a case or not, one can present it descriptively to oneself, for example, the good case as 'the good case' and the bad case as 'the bad case'. Since we have two cases and two modes of presentation of each of them, we have the four possibilities in Table 1 to consider.

Possibility II does not arise, because a case can be presented indexically as 'my present case' only if one is in it; since one cannot be in both the good and bad cases simultaneously, one cannot be faced with the task of discriminating between them, each presented indexically as 'my

¹ For epistemic asymmetry in relation to scepticism see Williams 1978: 310–13, although Williams is confident of an asymmetry only when death, drugs, sleep, or the like incapacitate the subject from thinking rationally. Humberstone 1988 has a subtle discussion of obstacles to asymmetry. For further sources of epistemic asymmetry, see section 10.4 and Appendix 5.

TABLE 1. Presentation of Cases

Possibility	Presentation of good case	Presentation of bad case
X II	Indexical: 'my case'	Indexical: 'my case'
ID	Indexical: 'my case'	Descriptive: 'the bad case'
DI	Descriptive: 'the good case'	Indexical: 'my case'
✓ DD	Descriptive: 'the good case'	Descriptive: 'the bad case'

present case'. DD discrimination is trivial, for one is merely required to discriminate conceptually between them presented as 'the good case' and 'the bad case', with no need to discover which case one is in. The interesting possibilities are ID and DI. Sceptics and anti-sceptics agree that in the bad case one cannot discriminate the bad case, presented indexically as 'my present case', from the good case, presented descriptively as 'the good case'. Thus it is uncontentious that the cases are DI indistinguishable. The issue is whether they are ID indistinguishable.

Indistinguishability is symmetric in the sense that if x presented under mode M is indistinguishable from y presented under mode N , then y presented under mode N is indistinguishable from x presented under mode M , but it obviously does not follow that x presented under mode N is indistinguishable from y presented under mode M . DI indistinguishability does not imply ID indistinguishability. The anti-sceptic claims that in the good case one can discriminate the good case, presented indexically as 'my present case', from the bad case, presented descriptively as 'the bad case', for that is just to know in the good case that one is not in the bad case. The sceptic claims that one cannot make that discrimination, but since that is in effect to claim that in the good case one cannot know that one is not in the bad case, ID indistinguishability is tantamount to the sceptic's conclusion. The sceptic cannot use it as a premise without begging the question.

In a more complex version of the argument, the sceptic may postulate a subject whose case oscillates over time between the good case and the bad case. Such a subject may indeed be incapable of discriminating between the good case, presented indexically as 'my present case', and the bad case, presented indexically as 'my case five minutes ago', and therefore lack the relevant knowledge. It does not follow that one lacks that knowledge even if one's case is not in fact oscillating, or in danger of doing so. Thus the oscillation example does not achieve the sceptic's purpose. Alternatively, the sceptic may prefer to work with identity of appearance rather than with indistinguishability. The ultimate uselessness of such an appeal will emerge in the course of the argument below.

8.3 DIFFERENCE OF EVIDENCE IN GOOD AND BAD CASES

The sceptic typically insists that one has exactly the same evidence in the two cases. Therefore, since one believes p with that evidence in the bad case, believing p with the evidence one has in the good case is insufficient for the truth of p . If the sceptic allowed that one had different evidence in the two cases, false belief in the bad case would be a far less pressing threat to knowledge in the good case: the possibility of falsely believing p on the basis of bad evidence is quite compatible with the possibility of knowing p on the basis of good evidence. Scepticism about the external world has more intuitive force than scepticism about one's own sensations because we do not usually envisage beliefs about one's own sensations as based on evidence insufficient for their truth.

The sceptic cannot simply stipulate that one has the same evidence in the good and bad cases. For the notion of evidence will serve the sceptic's purposes only if it has non-trivial connections with other epistemic notions, such as the notion of knowledge. Some externalists about evidence (although not all) will argue that those connections force a difference in evidence between the two cases. If the sceptic tries to stipulate that the bad case is a case in which one falsely believes p while having the same evidence as one has in a case in which by externalist standards one knows p , those externalists will reply that, so defined, the bad case is impossible, and the sceptic's argument does not get off the ground. Rather, the sceptic should define the bad case in less contested terms, so that its possibility is agreed, and then argue for the lemma that one has the same evidence in it as in the good case. Many contemporary non-sceptics accept that lemma in the sceptic's overall argument. They concede that when we have empirical knowledge, we could have had false belief in the same proposition with exactly the same evidence. Many hold that, at least in some contexts, the bad case is in some sense irrelevant to the attribution of knowledge in the good case.² For present purposes, what matters is simply the claim that one has the same evidence in the two cases. How can that claim be supported?

A natural argument is by *reductio ad absurdum*. Suppose that one has different evidence in the two cases. Then one can deduce in the bad case that one is not in the good case, because one's evidence is not what

² Lewis 1996 gives a recent account of this kind in which sameness of evidence plays a central role. McDowell 1982 denies that the evidence is the same. For the relevant alternatives approach generally see Goldman 1976, Stine 1976, Dretske 1981b, and Cohen 1988.

it would be if one were in the good case. But even the sceptic's opponent agrees that it is consistent with everything one knows in the bad case that one is in the good case. Therefore, one has the same evidence in the two cases.

The argument assumes that in the bad case one knows what one's evidence is, otherwise one would lack a premise for the deduction. Now, surely one can be rational even in the bad case; misleading evidence sometimes makes false beliefs rational. So one can know what one's evidence is, granted the assumption that rational thinkers are in a position to know what their evidence is. The appeal of that assumption is by no means limited to sceptics; after all, it says that rational thinkers are in a position to know something. The idea, already mentioned, is that rationality requires one to respect one's evidence, which one cannot expect to do without knowing what it is.

8.4 AN ARGUMENT FOR SAMENESS OF EVIDENCE

Let us analyse the argument for sameness of evidence in detail. For simplicity, we may concentrate on cases in which one is rational, possesses all the relevant concepts, and is currently reflecting on one's evidence and its implications; one is epistemically active enough to know what ever one is in a position to know about one's evidence. If the sceptic can show that under these conditions one's evidence is the same in the two cases, the anti-sceptic would have little to gain by insisting that it is different when one is less epistemically active.

We start with the premise that one knows what one's evidence is. 'Evidence' here and throughout means one's total body of evidence. To know what one's evidence is in the relevant sense, one must do better than merely to think of it as 'my evidence'. To be in a position to respect one's evidence, one must identify its specific content in a more perspicuous and intrinsic way. One need not compress the identification into a single item of knowledge. The content can be specified by a class of appropriate properties, each of which one knows one's evidence to have under some canonical specification of the property. We assume on behalf of the sceptic that for each appropriate property a unique canonical specification is given. Let us concede for the sake of argument that such a notion of the canonical can be worked out in detail. We may also assume that if a property is appropriate, so is its complement. The first premise is therefore:

- (1) For any appropriate property π , in any case in which one's evidence has π , one knows that one's evidence has π .

If we wanted to generalize (1) and the rest of the argument beyond cases in which one is rational, possesses all the relevant concepts, and is currently reflecting on one's evidence and its implications, we could replace 'knows' by 'is in a position to know'. One might wonder whether (1) generates an infinite regress, as Richard Fumerton (2000) has suggested. It does if being known to have π counts as an appropriate property whenever π does, but defenders of (1) should not concede that assumption. The appropriate properties are intrinsic to the content of one's evidence; being known to have such a property need not itself be intrinsic to the content of the evidence.

Whereas the first premise concerns first-personal knowledge of one's own case, the second concerns third-personal knowledge of one case from within another. For the argument to work, in the bad case one must know what one's evidence would be if one were in the good case, where the good case is presented descriptively. We can quite fairly assume that the terms 'the good case' and 'the bad case' abbreviate descriptions in which, for each appropriate property, if one's evidence in a case has an appropriate property then that is specified in the description of the case; likewise if one's evidence lacks an appropriate property. For the sceptic will insist that however much information one has about what would be so if one were in a given case, that still does not enable one to work out which case one is in. We may assume that one can refer to the appropriate properties, for that is already implicit in (1): if one's evidence has the appropriate property π , then one knows that it has π and so can refer to π ; if it lacks π , then it has the appropriate complementary property not- π , so one knows that it has not- π , so one can refer to not- π , so one can refer to π . Thus one can attain trivial conceptual knowledge in the bad case about the appropriate properties of one's evidence in the good case simply by unpacking one's descriptive concept of 'the good case':

- (2) For any appropriate property π , if in the good case one's evidence lacks π , then in the bad case one knows that in the good case one's evidence lacks π .

The third premise articulates the badness of one's predicament in the bad case. From premises each of which one knows in the bad case, one cannot deduce that one is not in the good case.

- (3) It is consistent with what one knows in the bad case that one is in the good case.

Now restrict ' π ' to appropriate properties and assume:

- (4) In the bad case one's evidence has π .

Suppose further, as an assumption for reductio ad absurdum:

- (5) In the good case one's evidence lacks π .

Premises (2) and (5) entail:

- (6) In the bad case one knows that in the good case one's evidence lacks π .

Premises (1) and (4) entail:

- (7) In the bad case one knows that one's evidence has π .

From 'In the good case one's evidence lacks π ' and 'One's evidence has π ' one can deduce 'One is not in the good case'. By (6) and (7), in the bad case one knows each premise of that deduction; hence:

- (8) It is inconsistent with what one knows in the bad case that one is in the good case.

Now (8), which rests on assumptions (1), (2), (4), and (5), contradicts (3). Thus on assumptions (1)–(4) we can deny (5) by reductio ad absurdum:

- (9) In the good case one's evidence has π .

We can conditionalize (9) on assumption (4):

- (10) If in the bad case one's evidence has π , then in the good case one's evidence has π .

Here (10) rests on assumptions (1)–(3). Since the appropriate properties were assumed to be closed under complementation, we can run through the argument (1)–(10) with 'not- π ' in place of ' π ', yielding:

- (11) If in the bad case one's evidence has not- π , then in the good case one's evidence has not- π .

Contraposition on (11) yields the converse of (10). Therefore, generalizing on ' π ' in (10) and (11), we have:

- (12) One's evidence in the good case has the same appropriate properties as one's evidence in the bad case.

The conclusion (12) rests on assumptions (1), (2), and (3). It may be restated as the claim that one's evidence is the same in the good and bad cases, where evidence is individuated by the appropriate properties. If

something like this argument is not the reason for which sceptics and others think that one has the same evidence in the two cases, it is not at all clear what is.

8.5 THE PHENOMENAL CONCEPTION OF EVIDENCE

That one has the same evidence in the good and bad cases is a severe constraint on the nature of evidence. It is inconsistent with the view that evidence consists of true propositions like those standardly offered as evidence for scientific theories. For example, the good case in which I see that the dial reads 0.407 corresponds to a bad case in which the dial does not read 0.407 but I hallucinate and it is consistent with everything I know that the dial reads 0.407. Since the proposition that the dial read 0.407 is false in the bad case, it is not evidence in the bad case. If my evidence is the same in the two cases, then that the dial read 0.407 is not evidence in the good case either. For similar reasons, (12) does not permit my evidence to include perceptual states individuated in part by relations to the environment. No matter how favourable my epistemic circumstances, I am counted as having only as much evidence as I have in the corresponding sceptical scenarios, no matter how distant and bizarre. Retinal stimulations and brain states fare no better as evidence, for in some sceptical scenarios they are unknowably different too. Thus (12) drives evidence towards the purely phenomenal.

We should not assume ourselves to grasp the concept of the phenomenal quite independently of (12). Instead, the phenomenal may be postulated as comprising those conditions, whatever they are, which rational subjects can know themselves to be in whenever they are in them. Such conditions may be supposed to comprise conditions on present memory experience as well as on present perceptual experience (Lewis 1996: 553). That such conditions exist is supposedly guaranteed by the argument that rationality requires one to respect one's evidence and cannot require one to respect something unless one is in a position to know what it is.³

³ Fumerton 2000 suggests an alternative conception of the phenomenal as that which supervenes on relations of direct acquaintance. Presumably, to be directly acquainted with something is to be acquainted with it but not by being acquainted with something else. But then we lack a sound argument to show that I cannot be directly acquainted with something in the good case—such as my hand—with which I am not directly acquainted in the bad case. Some of Fumerton's remarks also presuppose the equivalence of the notion of the phenomenal as what one is always in a position to know about with

The argument for (12) is not vulnerable to a distinction between relevant and irrelevant alternatives to the good case, for it in no way assumes the relevance of the bad case to the good case. It does not use the sceptical claim that it is consistent with what one knows in the good case that one is in the bad case; it uses only the uncontested claim (3) that it is consistent with what one knows in the bad case that one is in the good case. Although that may be to assume the relevance in some sense of the good case to the bad case, that assumption is uncontroversial, since the good case is the sort of case one believes oneself to be in and appears to oneself to be in if one is in the bad case. Even if in the good case one properly ignores the bad case, the argument to (12) still shows (given its premises) that one's evidence in the good case cannot exceed one's evidence in the bad case.

Does a distinction between relevant and irrelevant alternatives make trouble for the sceptic's further claim that false belief in the bad case precludes knowledge in the good case? Perhaps falsely believing p with given evidence in a case β precludes knowing p with the same evidence in a case α only if β is a relevant alternative to α in some sense of 'relevant' in which the bad case is not a relevant alternative to the good case. (Although that is not the present issue, it is difficult not to feel sympathy for the sceptic here. If one's evidence is insufficient for the truth of one's belief, in the sense that one could falsely believe p with the very same total evidence, then one seems to know p in at best a stretched and weakened sense of 'know'. We might contrast it with a more robust sense in which one knows the evidence itself, if evidence can be conceived propositionally. But all these questions presuppose that one's evidence is indeed the same in the good and bad cases. How compelling is the argument for (12)? In particular, how compelling is the justification of its crucial premise (1)?

8.6 SAMENESS OF EVIDENCE AND THE SORITES

We can undermine the argument for (12), and in particular its crucial premise (1), by constructing a parallel argument from (1) to a clearly false conclusion. Whatever the nature of evidence, rational thinkers do not always know what their evidence is. The argument exploits ordin-

a notion of the phenomenal as what one is infallible about. These notions are not obviously equivalent. I could be in a position to know p while falsely believing $\neg p$, because my guru tells me $\neg p$. If one can have contradictory beliefs, I might even know p while deceiving myself into believing $\neg p$.

ary limits to one's powers of discrimination. It is an application of the anti-luminosity argument of section 4.3, with some modifications to clarify its relation to the argument presented in section 8.4. The argument shows that the condition that one's evidence has the appropriate property π is not luminous; it can obtain even when one is not in a position to know that it obtains.

Let $t_0, t_1, t_2, \dots, t_n$ be a long sequence of times at one-millisecond intervals. Imagine that one's experience very gradually changes from t_0 to t_n ; for example, one watches the sun slowly rise. One loses exact track of time. One's evidence at the beginning of the process (pitch darkness) is quite different from one's evidence at the end (bright daylight). Some of the appropriate properties of one's evidence are different; for purposes of this argument, it does not matter whether the appropriate properties exhaust the content of one's evidence. We may assume that the complement of an appropriate property is itself an appropriate property, although the purpose of the argument could be achieved without that assumption. For $0 \leq i \leq n$, let α_i abbreviate a description of the case one is in at t_i ; the description specifies the time t_i in clock terms and lists the appropriate properties which one's evidence then has and those which it then lacks. As with the sceptic's original argument, we may assume that one can refer to the appropriate properties, for that is implicit in (1). Thus one can attain trivial conceptual knowledge in one case about the appropriate properties of one's evidence in another case simply by unpacking one's descriptive concept of the latter case; in particular:

- (2.) For any appropriate property π , if in α_{i-1} one's evidence lacks π , then in α_i one knows that in α_{i-1} one's evidence lacks π .

The justification of (2.) is just like the justification of (2.) above.

Now consider the description of what is in fact the case one was in a millisecond ago. Given one's limited powers of discrimination, one does not know propositions from which one can deduce that that description does not apply to one's own case:

- (3.) It is consistent with what one knows in α_i that one is in α_{i-1} .

Since the purposes of this chapter require only one example in which (1) has false consequences, any readers lucky enough to have perfect discrimination amongst their own states should consider the less fortunate example of the present author, who is frequently in a predicament like (3.). In such cases, (3.) is obvious in roughly the way in which it is obvious that it is consistent with what I know by sight when I am in fact looking at a distant tree i millimetres high that I am looking at a tree

only $i-1$ millimetres high. From premises which I know on the basis of sight to the conclusion that I am not looking at a tree only $i-1$ millimetres high, there is no hope of constructing a valid deduction, not even one which I am somehow not in a position to carry out. Similarly, from premises which I know in α , to the conclusion that I am not in α_{i-1} , there is no hope of constructing a valid deduction, not even one which I am somehow not in a position to carry out.

The argument proceeds as before. Restrict ' π ' to appropriate properties and assume:

- (4.) In α , one's evidence has π .

Suppose further, as an assumption for reductio ad absurdum:

- (5.) In α_{i-1} one's evidence lacks π .

Premises (2.) and (5.) entail:

- (6.) In α , one knows that in α_{i-1} one's evidence lacks π .

Premises (1) and (4.) entail:

- (7.) In α , one knows that one's evidence has π .

From 'In α_{i-1} one's evidence lacks π ' and 'One's evidence has π ' one can deduce 'One is not in α_{i-1} '. By (6.) and (7.), in α , one knows each premise of that deduction; hence:

- (8.) It is inconsistent with what one knows in α , that one is in α_{i-1} .

Now (8.), which rests on assumptions (1), (2.), (4.), and (5.), contradicts (3.). Thus on assumptions (1) and (2.)-(4.) we can deny (5.) by reductio ad absurdum:

- (9.) In α_{i-1} one's evidence has π .

We can conditionalize (9.) on assumption (4.):

- (10.) If in α , one's evidence has π , then in α_{i-1} one's evidence has π .

Here (10.) rests on assumptions (1), (2.), and (3.). Since the appropriate properties were assumed to be closed under complementation, we can run through the argument (1)-(10.) with 'not- π ' in place of ' π ', yielding:

- (11.) If in α , one's evidence has not- π , then in α_{i-1} one's evidence has not- π .

Contraposition on (11.) yields the converse of (12.). Thus, generalizing on ' π ' in (10.) and (11.), we have:

- (12.) One's evidence in α_{i-1} has the same appropriate properties as one's evidence in α .

Proposition (12.) rests on assumptions (1), (2.), and (3.). But the relation between the cases in (12.) is transitive; if one's evidence in case β has the same appropriate properties as one's evidence in case γ and one's evidence in case γ has the same appropriate properties as one's evidence in case δ , then one's evidence in β has the same appropriate properties as one's evidence in δ , for what is in question is exact sameness in all properties from a fixed class. Although (3.) claims only that α_{i-1} and α are indiscriminable, and indiscriminability is a non-transitive relation, we have deduced from it and the other premises the transitive relation of exact sameness of evidence in the appropriate respects. Thus (12.), ..., (12_n) together yield:

- (13) One's evidence in α_0 has the same appropriate properties as one's evidence in α .

The conclusion (13) rests on assumptions (1), (2₁), ..., (2_n), (3₁), ..., (3_n). But (13) is obviously false. One's evidence at the end of the process is grossly different from one's evidence at the beginning; it differs in many of its appropriate properties. Since (2₁), ..., (2_n), (3₁), ..., (3_n) are true, for reasons already given, (1) is false.

Even if we drop the assumption that the complements of appropriate properties are themselves appropriate, we still have the argument to (10.), and therefore by transitivity to the conclusion that if in α , one's evidence has an appropriate property, then in α_0 one's evidence already had that property. That is obviously false, too. One does not always know the appropriate properties of one's evidence; one does not always know what one's evidence is.

To the objection that the argument is undermined by its obvious similarity to a sorites paradox, the reply is just as in section 4.5, and will not be repeated here. In brief, the argument in a sorites paradox has an obviously false premise when the vague terms at issue are sharpened, here that is not so.

Fumerton (2000) points out that the sorites argument would show that (1) can fail in a small way; it would not show that (1) can fail in a large way, as is held to occur in the bad case. However, one's evidence in the bad case can appear exactly similar to one's evidence in the good case, not because it is almost exactly similar, but because it is so radically impoverished that one lacks evidence of its impoverishment. Moreover, the usual reasons for claiming that one is always in a position to know exactly what one's evidence is do not naturally evolve into reasons for claiming that one is always in a position to know approximately what one's evidence is. We are often in a position to know approximately what our evidence is; that our position

should occasionally be much worse than that, as in the bad case, is no surprise.⁴

8.7 THE NON-TRANSPARENCY OF RATIONALITY

The argument against (1) does not depend on any specific theory of evidence. The crucial assumption about evidence is just that its appropriate properties can vary between the endpoints of a spectrum of cases, as they must if we are to learn from experience. Whatever evidence is, one is not always in a position to know what one has of it. Thus nothing would be gained by a retreat to the fallback claim that one always knows (or is in a position to know) what one's evidence *appears* to be. For we can replace the words 'one's evidence has [lacks] π ' in the preceding argument by 'one's evidence appears to have [lack] π '. Under this modification, (1) expresses the fallback claim, $(2_1), \dots, (2_n)$ can be justified in the same way as before, $(3_1), \dots, (3_n)$ are unchanged, and (13) remains hopelessly implausible, so the argument refutes the fallback claim, too. One does not always know what one's evidence appears to be.

If the phenomenal is postulated as comprising those conditions of the subject, whatever they are, which are accessible to the subject whenever they obtain, and therefore satisfy something like desideratum (1) for evidence, then the phenomenal is empty. We have the illusion of coming ever closer to a phenomenal core of experience by progressively eliminating every feature which can fail to be accessible to the subject, but, like the sequence of open intervals $(0,1), (0,1/2), (0,1/4), \dots$, this sequence of approximations converges to the empty set.

We could modify (1) by relativizing appropriateness to cases. The modified variant of (1) would claim that, for any case α and any property π appropriate to α , if in α one's evidence has π , then one knows in α that one's evidence has π . We could then no longer argue to (13), because 'appropriate' in the modified (12_i) would have different relativizations for different values of i . But the argument for sameness of evidence in the good and bad cases would fail, for although we could

⁴ Following Poincaré, Russell used the non-transitivity of indistinguishability in sensation to argue for imperceptible differences amongst our sense data (Russell 1993: 148, originally published in 1914). A. J. Ayer replied that the only notion of exact resemblance applicable to sense data is equivalent to the relation of apparent exact resemblance between material things, which can be non-transitive (1940: 132-4). That reply provides no basis for resistance to the arguments of this chapter.

show that one's evidence in the good case had the same properties appropriate to the bad case as one's evidence in the bad case, we could not show that one's evidence in the good case had the same properties appropriate to the good case as one's evidence in the bad case. Indeed, we could not show that one was always in a position to know which properties of evidence were appropriate to one's own case. The proposed relativization plays into the hands of the present strategy.

The problem remains: how can rational thinkers respect their evidence if they do not know what it is? If rationality requires one to respect one's evidence, then it is irrational not to respect one's evidence. But how can failing to respect one's evidence be irrational when one is not in a position to know that one is failing to respect one's evidence? More generally, how can ϕ -ing be irrational when one is not in a position to know that one is ϕ -ing?

The standard conception of rationality depends on a distinction between the *aims* and *methods* of cognitive activity. On that conception, truth is an aim. We cannot attain it directly; we cannot follow the rule 'Believe truly!' when we do not know what is true. Therefore we must use methods to reach the truth. Rationality is a method. We can follow rules of rationality because we are always in a position to know what they require. If the argument of section 8.6 is correct, this picture of rationality is mistaken. Just as one cannot always know what one's evidence is, so one cannot always know what rationality requires of one. Just like evidence, the requirements of rationality can differ between indiscriminable situations. Rationality may be a matter of doing the best one can with what one has, but one cannot always know what one has, or whether one has done the best one can with it. If something is a method only if one is always in a position to know whether one is complying with it, then there are no methods for learning from experience. But that standard is too exacting to be useful. We can use something as a method in contexts in which one is usually in a position to know whether one is complying with it, even if in other contexts one is not usually in a position to know whether one is complying with it. In that sense, we can use even believing truly as a method in contexts in which one is usually in a position to know what is true: for example, when forming beliefs in normal conditions about the spatial arrangement of medium-sized objects in one's immediate environment. In more difficult contexts, believing truly becomes an aim and we fall back on the method of believing rationally. Rationality becomes a sub-goal on the way to truth. That does not require one always to be in a position to know what rationality requires of one; it requires merely that one often knows what rationality requires when one does not know what truth

requires. Nothing has been said here to undermine that requirement. In still more problematic contexts, paradoxes throw our very standards of rationality into doubt, and we fall back still further on what workable methods we can find. Cognition is irremediably opportunistic.

There is a pragmatist and subjective Bayesian project to operationalize epistemology by working only with concepts whose application is always accessible to the agent. The argument of this chapter implies that the project is doomed to failure.

Uncertainty about evidence does not generate an infinite regress of evidence about evidence about . . . In order to reflect adequately on one's evidence, one might need evidence about one's evidence, and in order to reflect adequately about the latter evidence, one might need evidence about it, and so on. But this regress is merely and harmlessly potential. We cannot in fact realize infinitely many levels of adequate reflection; at best, further reflection enables us to realize finitely many further stages. At some stage one must rely on unreflective causal sensitivity to evidence (see section 9.3).

One can be causally sensitive to a factor without being in a position to have exact knowledge of it, as when one is causally sensitive through unaided perception to the distances between objects in one's environment. One can be causally sensitive to appropriate properties of one's evidence without being in a position to know them exactly. Causal sensitivity need not be perfect to be genuine. Sufficiently bad cognitive circumstances may involve obstacles even to causal sensitivity to one's evidence. The bad case in a sceptical argument may be a case in point. One's cognitive circumstances may be so bad that one is in no position to know how impoverished one's evidence is in comparison to the good case. Our causal insensitivity to any difference in evidence between the two cases does not show that there is no difference in evidence between them.

It has not been shown that the good and bad cases do differ in evidence. That requires a positive account of evidence, which Chapters 9 and 10 will develop. They defend the view that one's total evidence (not one's evidence for p alone) is simply one's total knowledge, on which the assumption that one has the same total evidence in the two cases is tantamount to the sceptic's conclusion. For since, uncontroversially, in the bad case one fails to know p , p would not be part of one's total evidence in the bad case, and would therefore not be part of one's total evidence in the good case either; so in the good case, too, one would not know p . A sceptic who assumes that one's total evidence is the same begs the question against a non-sceptic who takes that view of evidence. Of course, the argument of this chapter does not assume the equation of

one's total evidence with one's total knowledge; rather, it lays the groundwork for the equation.⁵

For present purposes, what matters is that sameness of evidence has not been established, and a salient argument for it has turned out to rest on a false premise. For all that the sceptic has shown, one has more evidence in the good case than in the bad case, and knowledge in the former is unthreatened by false belief in the latter.

The problem is not confined to the sceptic. It also affects those non-sceptics who argue that in the good case we know p even though we falsely believe p in some cases in which we have the same evidence as in the good cases, because those bad cases are irrelevant (at least in this context). Such theorists have not eliminated the hypothesis that one knows p only if one does not falsely believe p on the same evidence in any case at all, relevant or irrelevant. That hypothesis does not entail scepticism.

Contextualists may argue that the extension of 'evidence' waxes and wanes with the context of utterance just as they suppose the extension of 'knowledge' to do. But then they cannot use 'the same evidence' as a fixed standard against which to measure contextual variation in standards of relevance. 'One knows p ' is supposed to count as true in the good case when the bad case is irrelevant; but if (speaking in such a context) the bad case also counts as differing from the good case in non-pragmatic respects such as evidence, why invoke pragmatic respects such as relevance?

8.8 SCEPTICISM WITHOUT SAMENESS OF EVIDENCE

Sometimes the good and bad cases in a sceptical argument have a different structure from that considered so far. Scepticism about p does not always require the (metaphysical) possibility of a bad case in which one falsely believes p . Let p be a mathematical truth, and therefore a necessary truth. Thus no case in which one falsely believes p is possible; yet one can still doubt p , by doubting the reliability of the methods which led one to believe p . After all, someone with great faith in a certain coin might decide to believe p if it comes up heads and to believe $\neg p$ if it

⁵ The equation allows sameness of evidence between two cases in which, without knowing p , one justifiably believes p , in one case truly, in the other falsely. Fumerton (2000) describes such a case, under the impression that it constitutes a difficulty for the equation.

comes up tails; if he believes p because the coin came up heads, he does not know p , although he could not have believed p falsely. His belief fails to be knowledge because the method by which he reached it could just as easily have led to a false belief in a different proposition (see also section 4.4). His evidence in the bad case includes the proposition that the coin came up tails, and therefore differs from his evidence in the good case. Even when false belief in p is possible, one's evidence in the bad case which motivates scepticism about p need not be the same as in the good case. Incoherent dreams feel coherent; one might therefore doubt the coherence of one's present experience, even though it feels, and in fact is, coherent. Since one's experience is coherent in the good case and incoherent in the bad case, one's evidence presumably differs between the two cases. The sceptic need not claim that in some possible case, one's experience is incoherent and one has the same evidence which one actually has, for that would be too close to asserting dogmatically that one's actual experience is incoherent. Rather, the locus of the doubt is the method by which one reaches the belief that one's experience is coherent.

Does the discussion of sceptical arguments in sections 8.2–8.7 generalize to these examples? One might think not. 'Method' has replaced 'evidence' as the crucial term, and they seem to be crucially disanalogous: we are far more strongly tempted to assume that one is always in a position to know what one's evidence is than to assume that one is always in a position to know what method one is using. The sceptic will happily allow that our beliefs may have inaccessible, unconscious causes, and argue that for all we know such causes are quite insensitive to whether the beliefs they cause are true. What if we are in fact using a method which cannot yield false beliefs? The sceptic will point to cases in which we merely appear to be using that method and our resulting beliefs are false. Such cases are supposed to falsify our knowledge claims. For the sceptic, the methods on whose reliability the epistemic status of our beliefs depends are individuated by appearances; the falsity of beliefs reached in cases which appear the same as the actual case in respect of one's method constitutes unreliability in one's actual apparent method.

In effect, the sceptic distinguishes the *process* by which one's belief was caused from the *rule* which one used in reaching the belief. Processes are at the subpersonal level, rules at the personal level. One can take responsibility for one's rules in a way in which one cannot for the processes. One's rationality depends on the rules which one uses rather than the processes which go on in one. One is typically not in a position to know what process caused one's belief, but we are tempted

to suppose that one is always in a position to know what rules one used in reaching it. For if one was not in a position to know what rules one was using, and one's rationality depended on the rationality of one's rules, how could one be required to be rational? The requirement that one is always in a position to know what rules one is using forces us into individuating them phenomenally, just as the corresponding epistemic requirement on evidence forced us into individuating it phenomenally. If a rule were individuated non-phenomenally, one could to all appearances be using it while not in fact be doing so; in which case one would not be in a position to know that one was not using that rule. An argument similar to (1)–(12) concludes that one is using the same rule in the good and bad cases. According to the sceptic, what matters for knowledge is the reliability of the rule, not of the process, so one's false belief in the bad case makes one's rule in the good case unreliable, and therefore undermines knowledge in the good case.

The sceptic's conception of a rule collapses. By an argument parallel to (1)–(13) (via (2₁)–(12₁)), only trivial rules meet the epistemic requirement. For a series of indiscriminable differences links a case in which one uses a given rule to a case in which one uses a quite different rule. For example, one initially believes p for reason R while giving no weight to reason R^* ; gradually one gives less weight to R and more to R^* , until finally one believes p for reason R^* while giving no weight to R . R and R^* differ so much in kind that believing for reason R and believing for reason R^* amount to using different rules. An argument just like that of section 8.6 refutes the assumption that in every case one is in a position to know what rule one is using. Even when the sceptic does not assume identity of evidence between the good and bad cases, the underlying dialectic is the same.

We can fall into scepticism if we attribute too much self-knowledge to the subject in bad cognitive circumstances, for the asymmetry in knowledge between the good and bad cases requires an asymmetry in self-knowledge. Once we relax our claims to self-knowledge, we strengthen our claim to knowledge of the external world. Sceptical arguments fail when they depend on exempting an internal world of appearances, for they depend on misconceiving appearances as just what they appear to be. The ruthless sceptic grants no exemptions. If the sceptic must argue that we never even know how things appear to us, should we still harbour the sneaking suspicion that scepticism is right after all?